

(12) **United States Patent**  
Navasivasakthivelsamy et al.

(10) **Patent No.:** US 10,652,020 B2  
(45) **Date of Patent:** May 12, 2020

(54) **SWITCHING DATA SOURCES DURING ROLLING UPGRADE OF A LEADER-BASED DISTRIBUTED SERVICE WITH MINIMAL DISRUPTIONS**

(71) Applicant: **Nutanix, Inc.**, San Jose, CA (US)

(72) Inventors: **Arun Navasivasakthivelsamy**, San Jose, CA (US); **Gregory A. Smith**, Seattle, WA (US)

(73) Assignee: **NUTANIX, INC.**, San Jose, CA (US)

(\*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 218 days.

(21) Appl. No.: **15/842,624**

(22) Filed: **Dec. 14, 2017**

(65) **Prior Publication Data**  
US 2018/0173573 A1 Jun. 21, 2018

**Related U.S. Application Data**  
(60) Provisional application No. 62/434,873, filed on Dec. 15, 2016.

(51) **Int. Cl.**  
**H04L 9/32** (2006.01)  
**H04L 9/14** (2006.01)  
(Continued)

(52) **U.S. Cl.**  
CPC ..... **H04L 9/32** (2013.01); **G06F 8/656** (2018.02); **G06F 9/45512** (2013.01);  
(Continued)

(58) **Field of Classification Search**  
CPC . H04L 9/32; H04L 9/14; H04L 9/0844; H04L 67/141; H04L 63/0428;  
(Continued)

(56) **References Cited**

U.S. PATENT DOCUMENTS

8,549,518 B1 10/2013 Aron et al.  
8,601,473 B1 \* 12/2013 Aron ..... G06F 9/45533  
718/1

(Continued)

OTHER PUBLICATIONS

Poitras, Steven. "The Nutanix Bible" (Jan. 11, 2014), from <http://stevenpoitras.com/the-nutanix-bible/> (Publication date based on indicated capture date by Archive.org; first publication date unknown).

(Continued)

*Primary Examiner* — Wynuel S Aquino

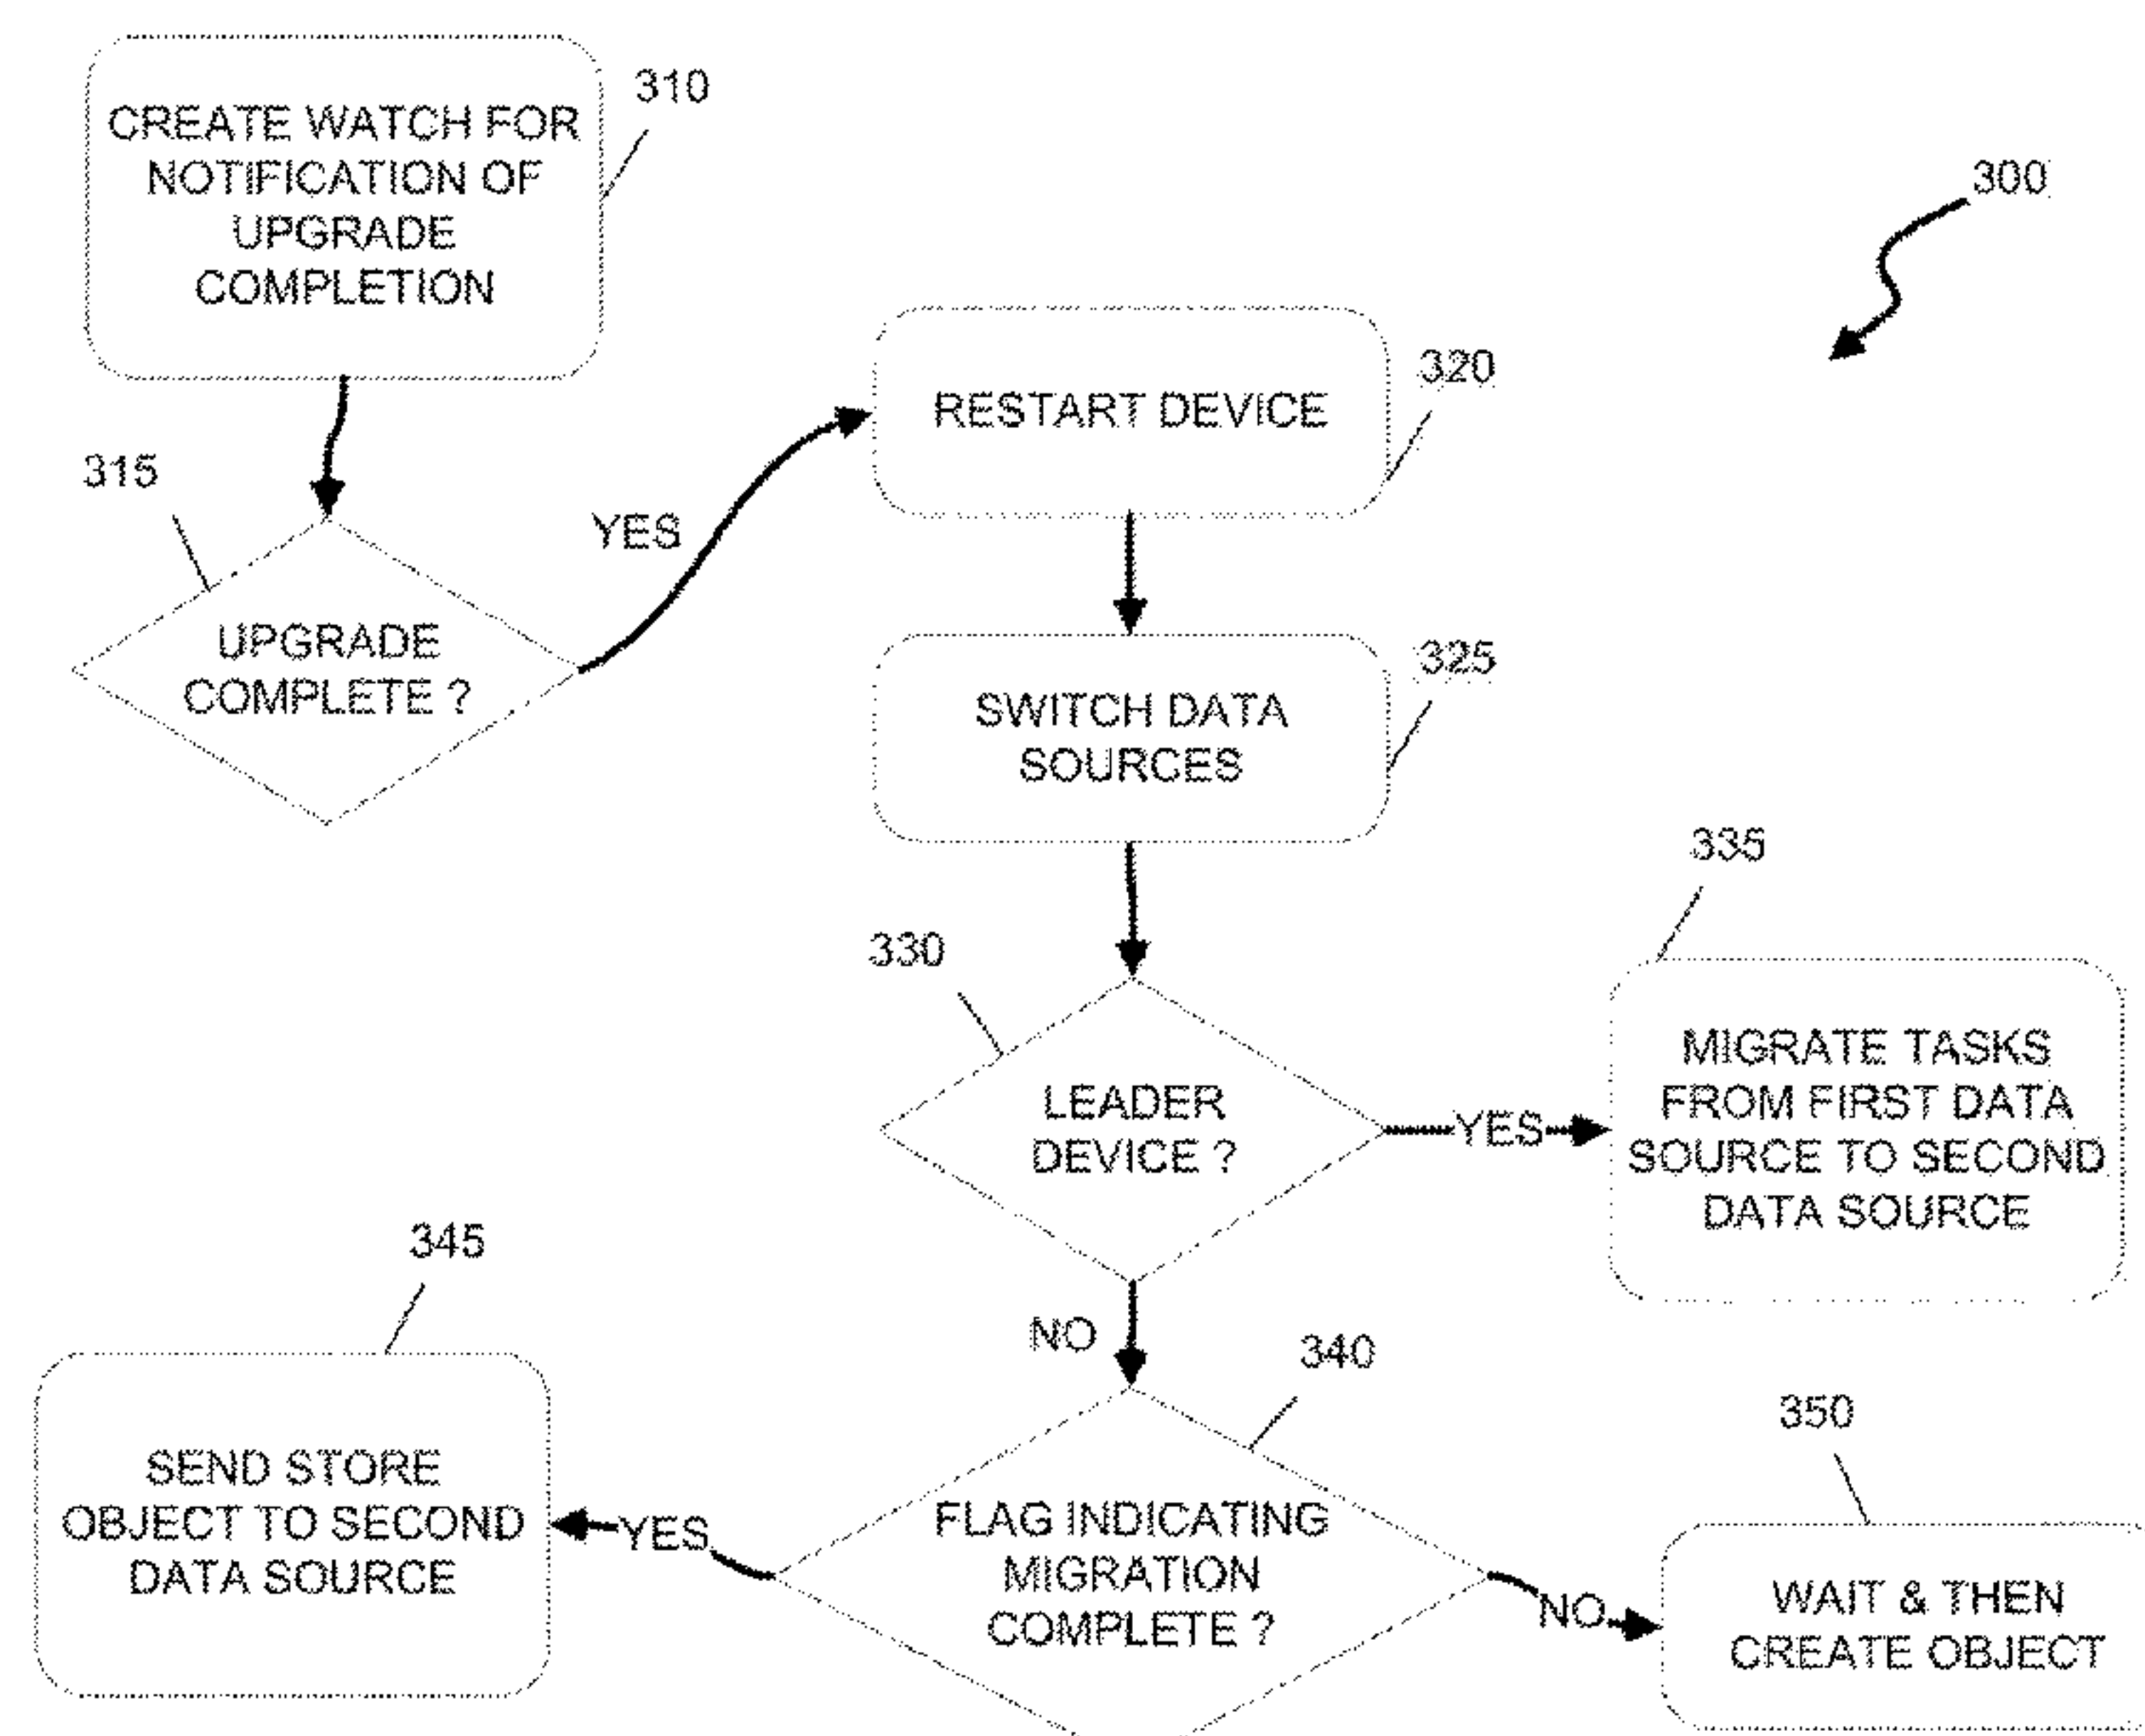
*Assistant Examiner* — Kevin X Lu

(74) *Attorney, Agent, or Firm* — Foley & Lardner LLP

(57) **ABSTRACT**

The disclosure relates to a method of switching data sources during a rolling upgrade of a leader-based distributed service. In at least one implementation, the method includes establishing, at a computing device in a cluster of devices in a network, a watch for notification of completion of a rolling upgrade on all devices in the cluster; restarting the computing device on completion of the rolling upgrade; switching, at the computing device, from a first data source to a second data source once the rolling upgrade is complete on all devices in the cluster; if the computing device is a leader device in the cluster, migrating tasks from the first data source to the second data source after completion of the rolling upgrade; and if the computing device is a non-leader device in the cluster, determining if task migration to the second data source is done and, when done, sending a store object to the second data source.

**21 Claims, 4 Drawing Sheets**



- (51) **Int. Cl.**  
*H04L 29/08* (2006.01)  
*G06F 9/50* (2006.01)  
*G06F 9/455* (2018.01)  
*G06F 21/44* (2013.01)  
*H04L 29/06* (2006.01)  
*H04L 9/08* (2006.01)  
*G06F 8/656* (2018.01)  
*G06F 9/46* (2006.01)
- (52) **U.S. Cl.**  
CPC ..... *G06F 9/45558* (2013.01); *G06F 9/461* (2013.01); *G06F 9/5077* (2013.01); *G06F 9/5088* (2013.01); *G06F 21/44* (2013.01); *H04L 9/0844* (2013.01); *H04L 9/14* (2013.01); *H04L 63/0428* (2013.01); *H04L 63/08* (2013.01); *H04L 67/141* (2013.01); *G06F 2009/4557* (2013.01); *G06F 2009/45591* (2013.01); *G06F 2209/505* (2013.01)
- (58) **Field of Classification Search**  
CPC ... H04L 63/08; G06F 9/5077; G06F 9/45512; G06F 9/45558; G06F 9/461; G06F 9/5088; G06F 21/44; G06F 8/656  
See application file for complete search history.

(56) **References Cited**  
U.S. PATENT DOCUMENTS

8,850,130	B1	9/2014	Aron et al.	
9,772,866	B1	9/2017	Aron et al.	
9,971,589	B2 *	5/2018	Chittigala	G06F 8/65
2006/0085824	A1 *	4/2006	Bruck	H04N 7/17336 725/86
2016/0112241	A1	4/2016	Kaushik	
2016/0197844	A1 *	7/2016	Smith	H04L 67/1004 709/226

2016/0203166	A1 *	7/2016	Aron	G06F 9/45558 707/609
2016/0246833	A1 *	8/2016	Eberlein	G06F 16/2315
2017/0364345	A1 *	12/2017	Fontoura	G06F 8/65
2018/0032323	A1 *	2/2018	Roy	G06F 8/65
2018/0165157	A1 *	6/2018	Olinsky	G06F 11/1433
2018/0176016	A1 *	6/2018	Deodhar	H04L 63/0428

OTHER PUBLICATIONS

Poitras, Steven. "The Nutanix Bible" (Jan. 12, 2016), from <https://nutanixbible.com/>.

Poitras, Steven. "The Nutanix Bible" (Oct. 15, 2013), from <http://stevenpoitras.com/the-nutanix-bible/> (Publication date based on indicated capture date by Archive.org; first publication date unknown).

Poitras, Steven. "The Nutanix Bible" (Jun. 20, 2014), from <http://stevenpoitras.com/the-nutanix-bible/> (Publication date based on indicated capture date by Archive.org; first publication date unknown).

Poitras, Steven. "The Nutanix Bible" (Jan. 3, 2017), from <https://nutanixbible.com/>.

Poitras, Steven. "The Nutanix Bible" (Sep. 4, 2015), from <https://nutanixbible.com/>.

Poitras, Steven. "The Nutanix Bible" (Jan. 7, 2015), from <http://stevenpoitras.com/the-nutanix-bible/> (Publication date based on indicated capture date by Archive.org; first publication date unknown).

Poitras, Steven. "The Nutanix Bible" (Jun. 8, 2017), from <https://nutanixbible.com/>.

Poitras, Steven. "The Nutanix Bible" (Jun. 9, 2015), from <http://stevenpoitras.com/the-nutanix-bible/> (Publication date based on indicated capture date by Archive.org; first publication date unknown).

Poitras, Steven. "The Nutanix Bible" (Jun. 9, 2016), from <https://nutanixbible.com/>.

Cano, Ignacio, et al. "Curator: Self-Managing Storage for Enterprise Clusters" (Mar. 27, 2017), from <https://www.usenix.org/conference/nsdi17/>.

\* cited by examiner



FIG. 1A

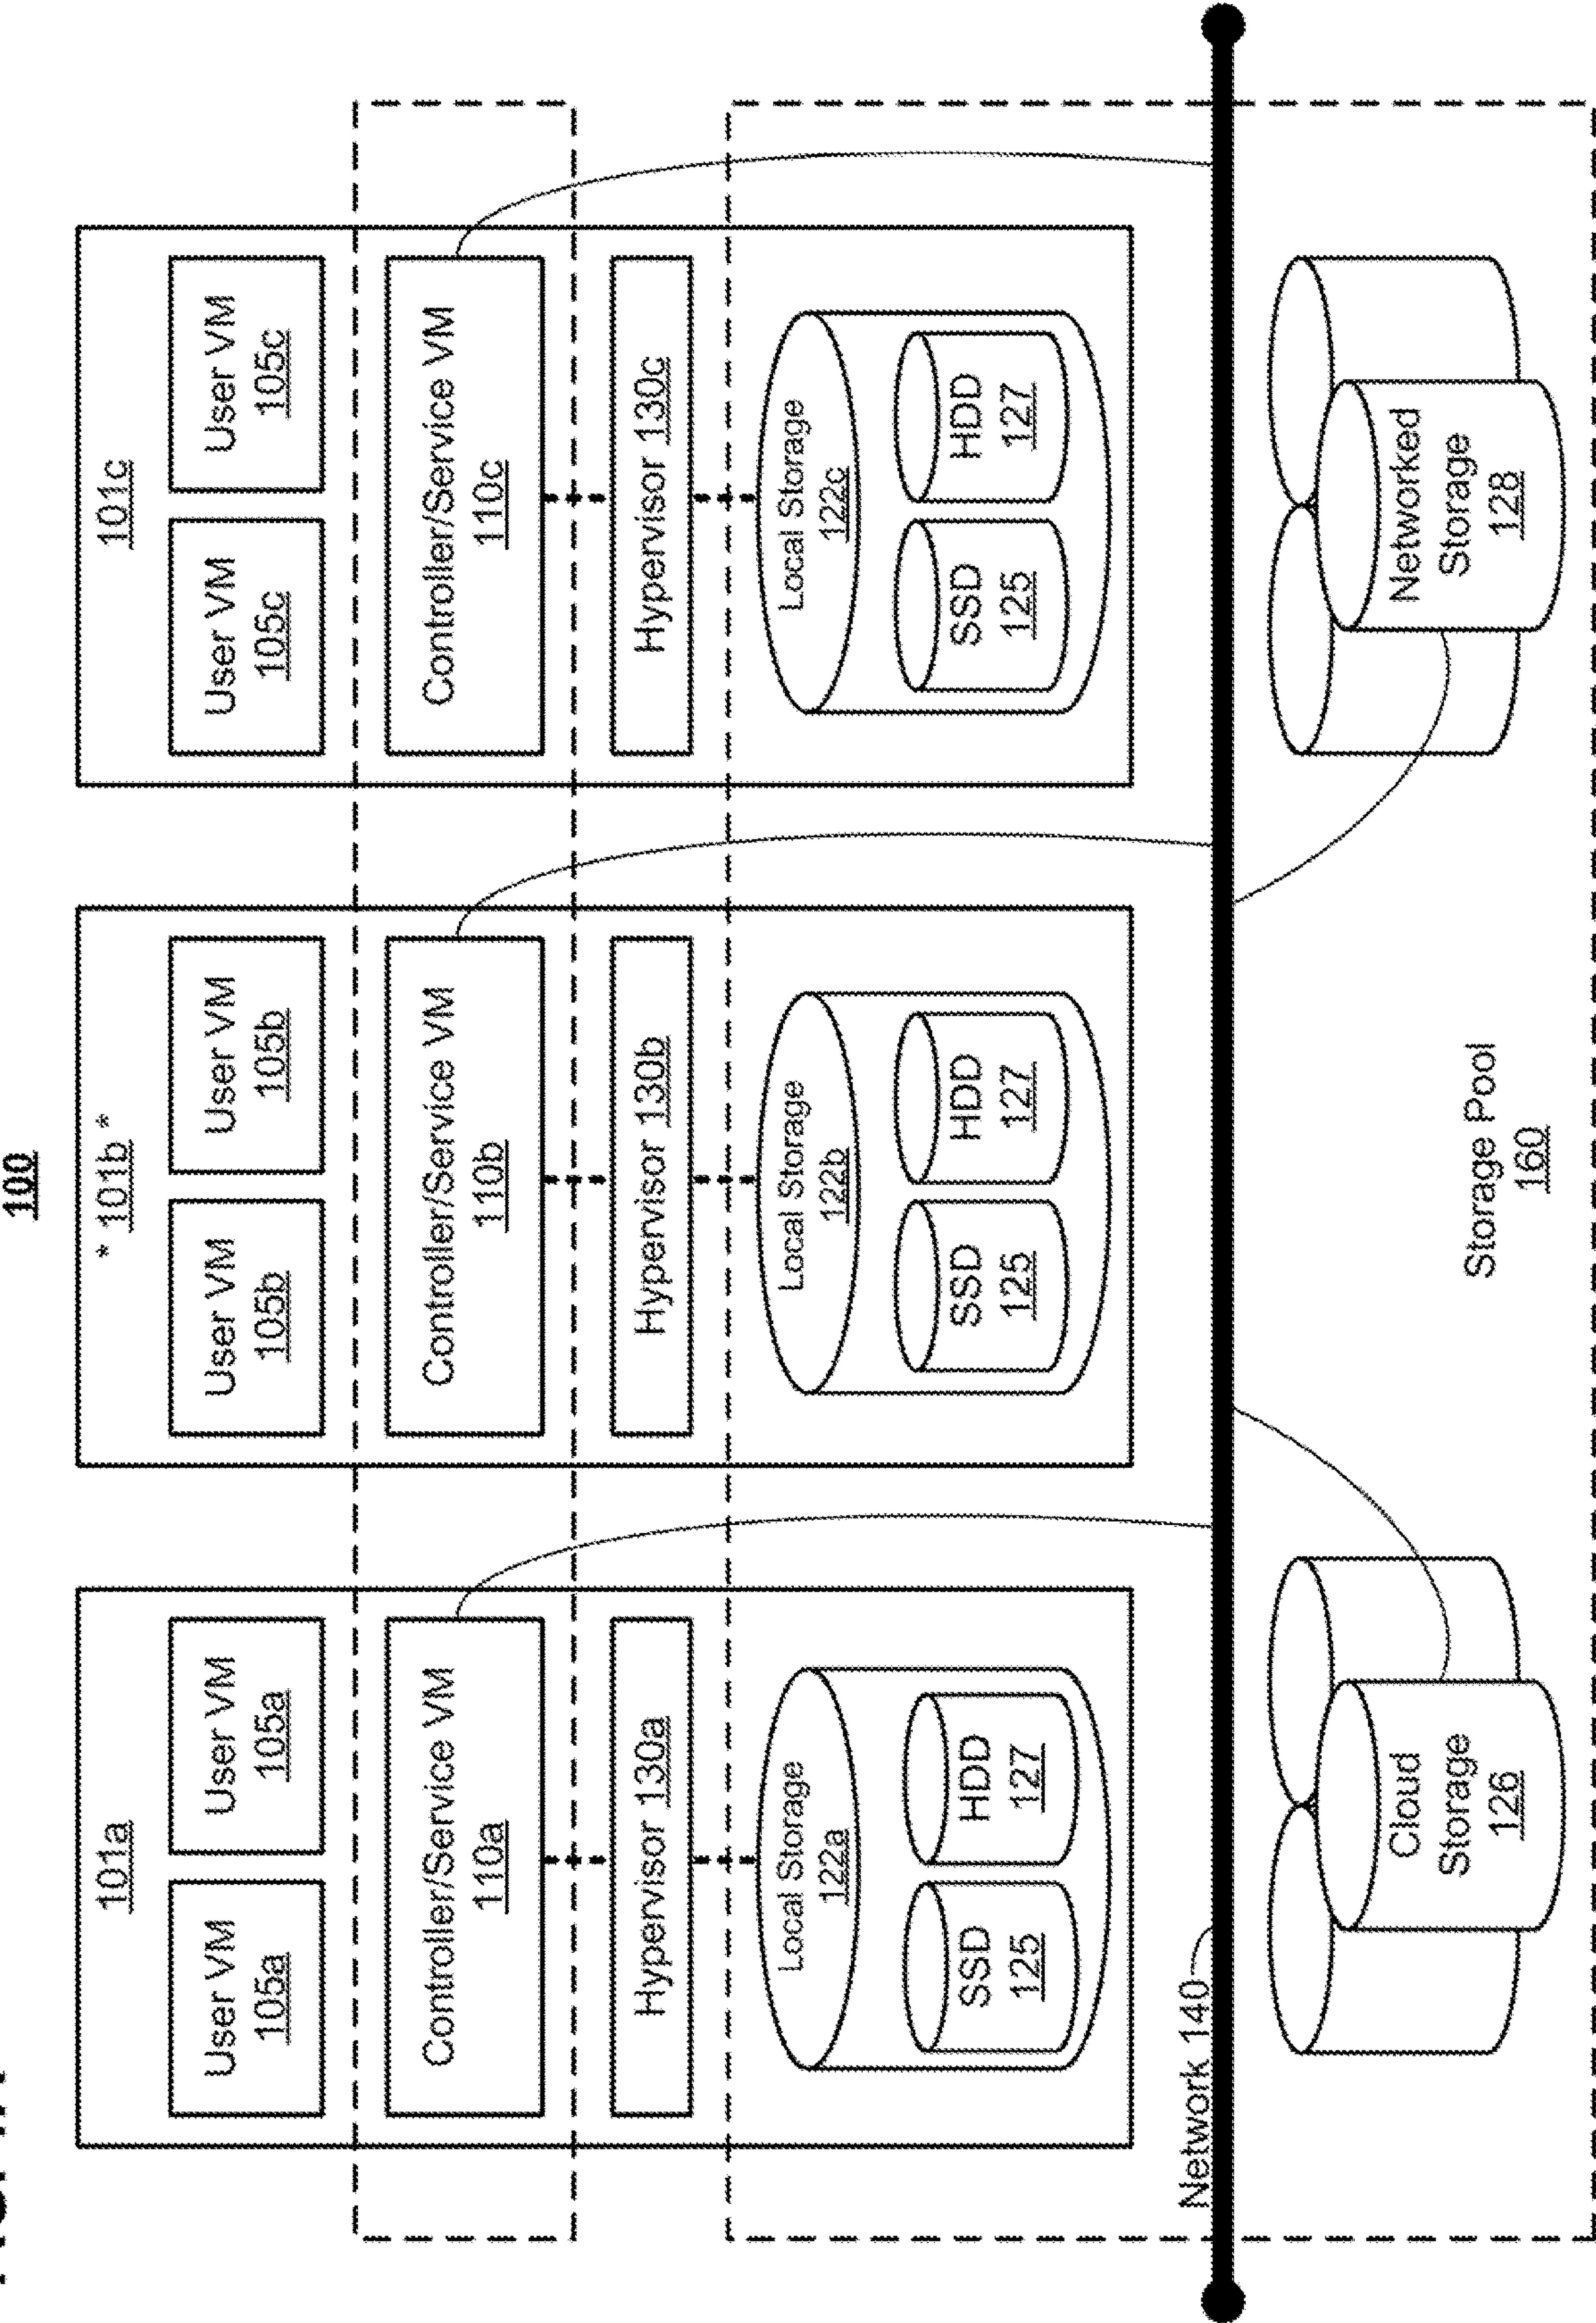




FIG. 1B

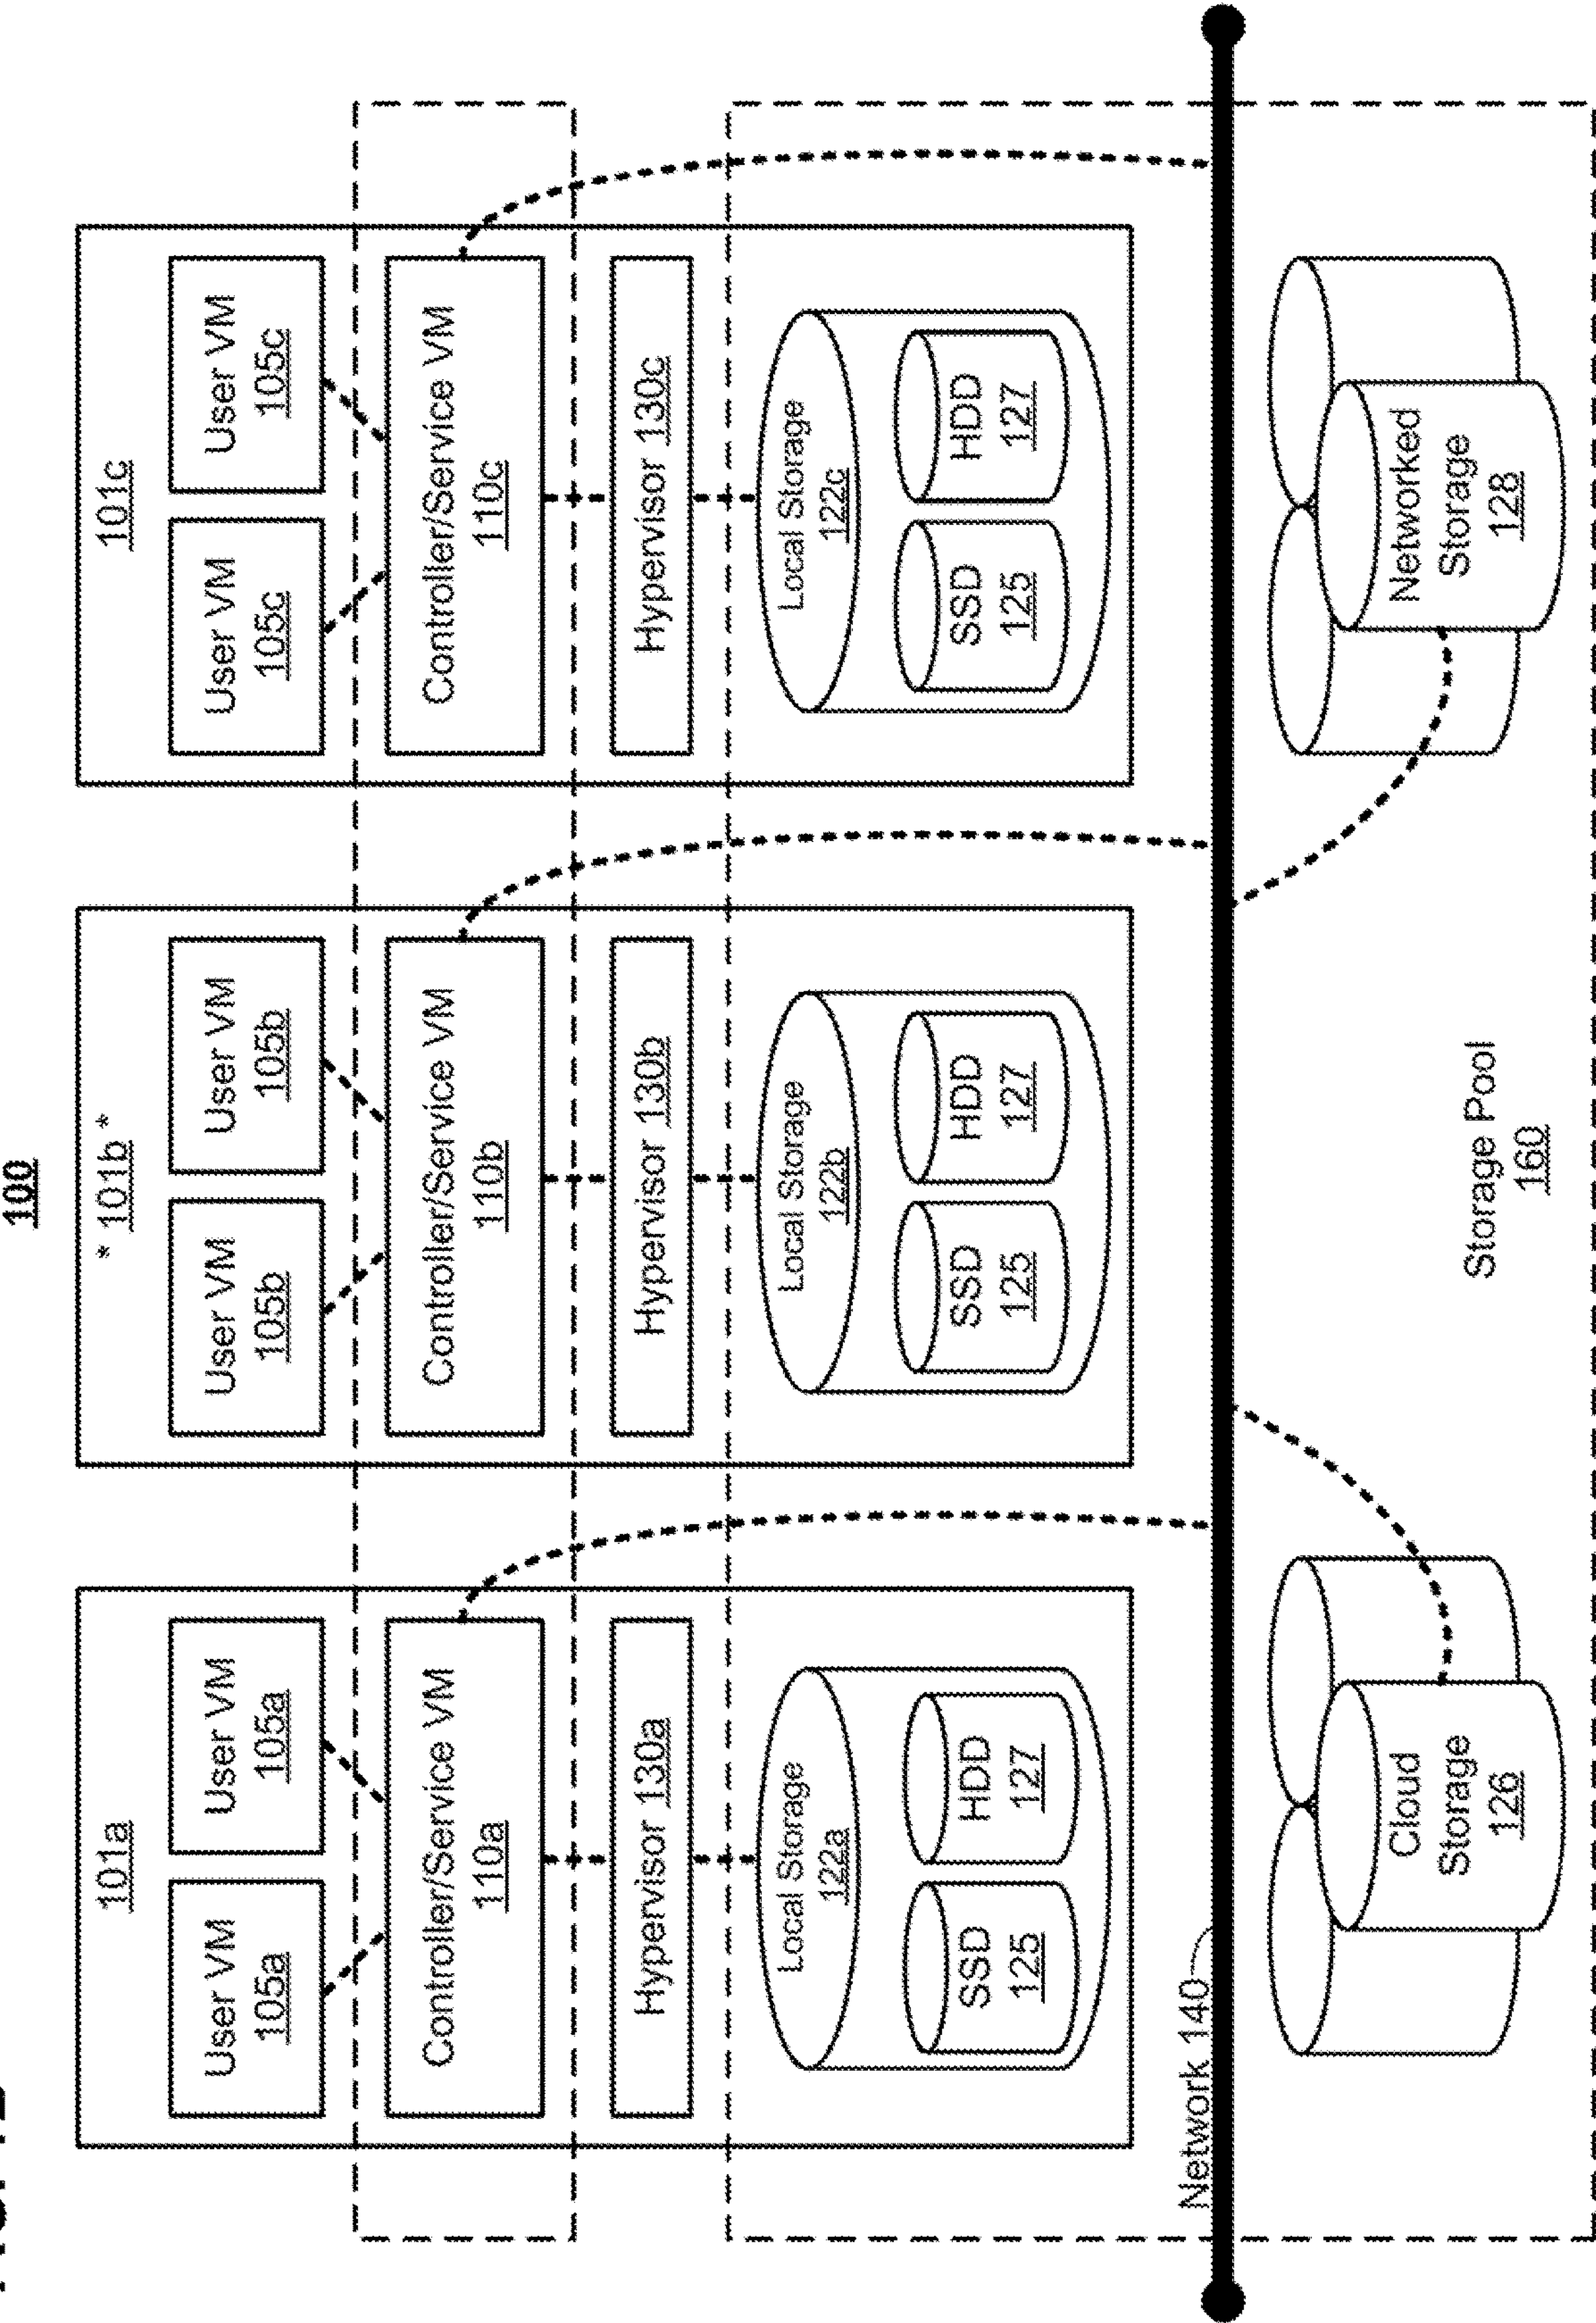
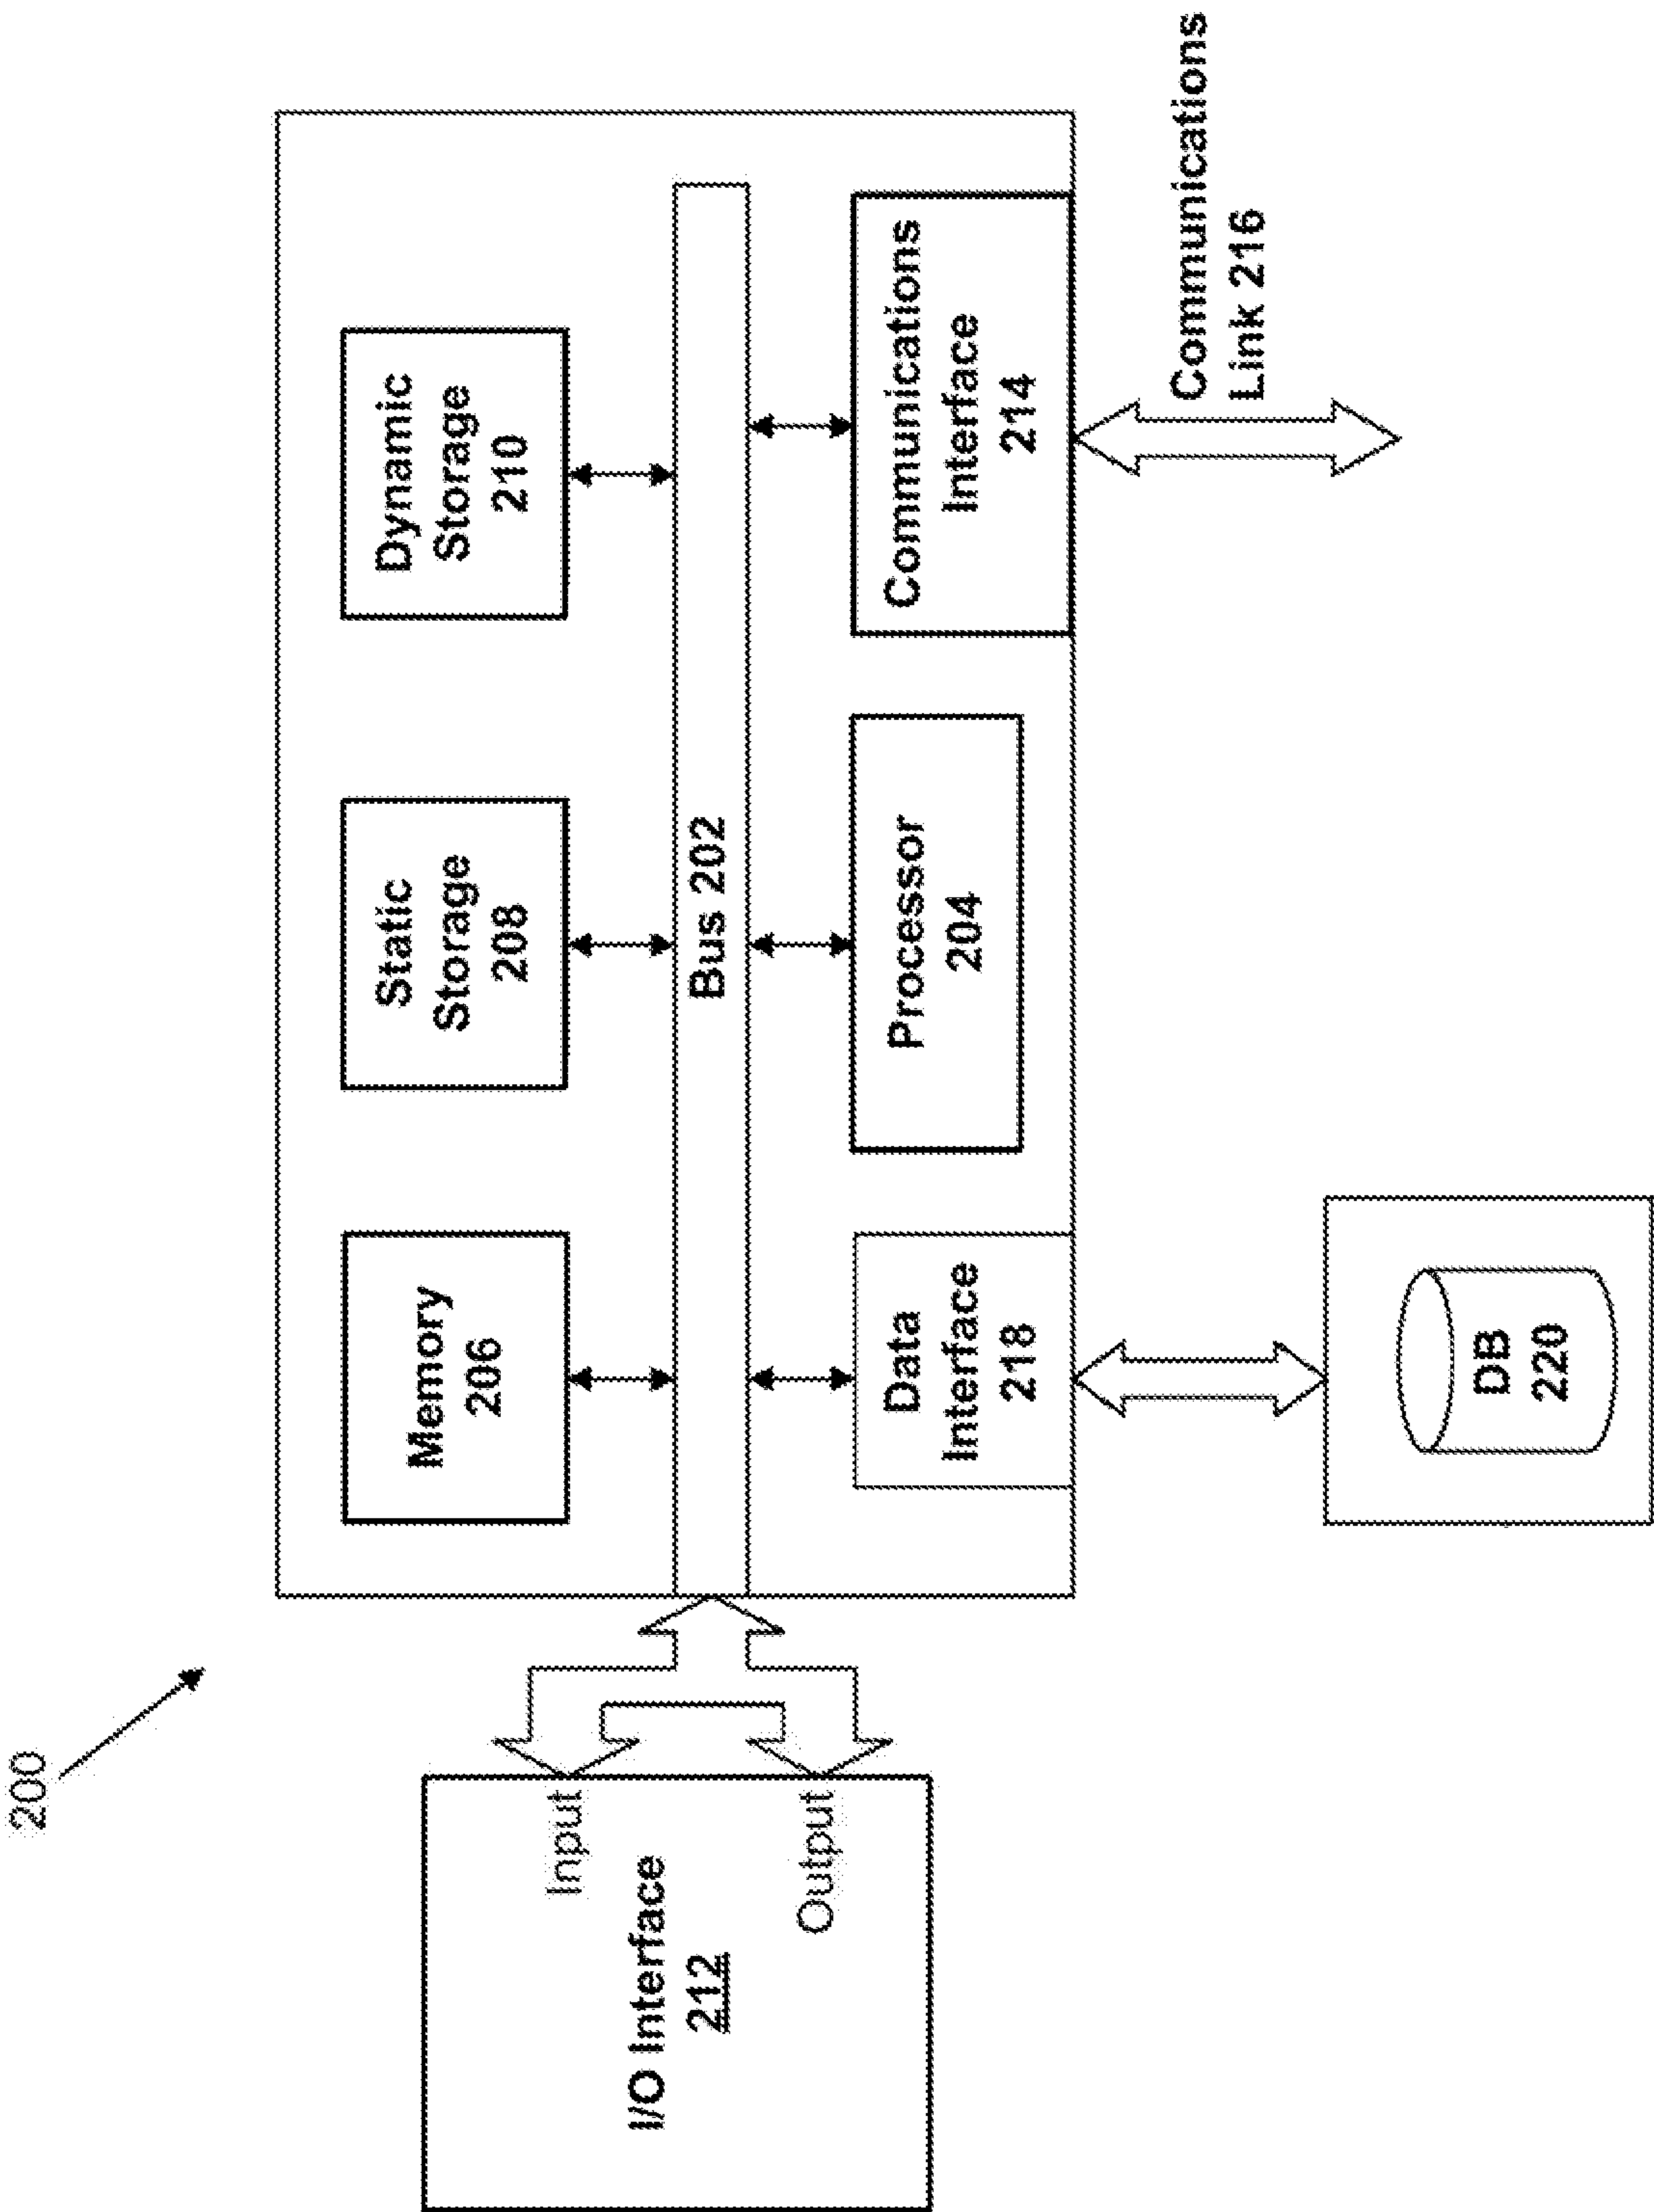
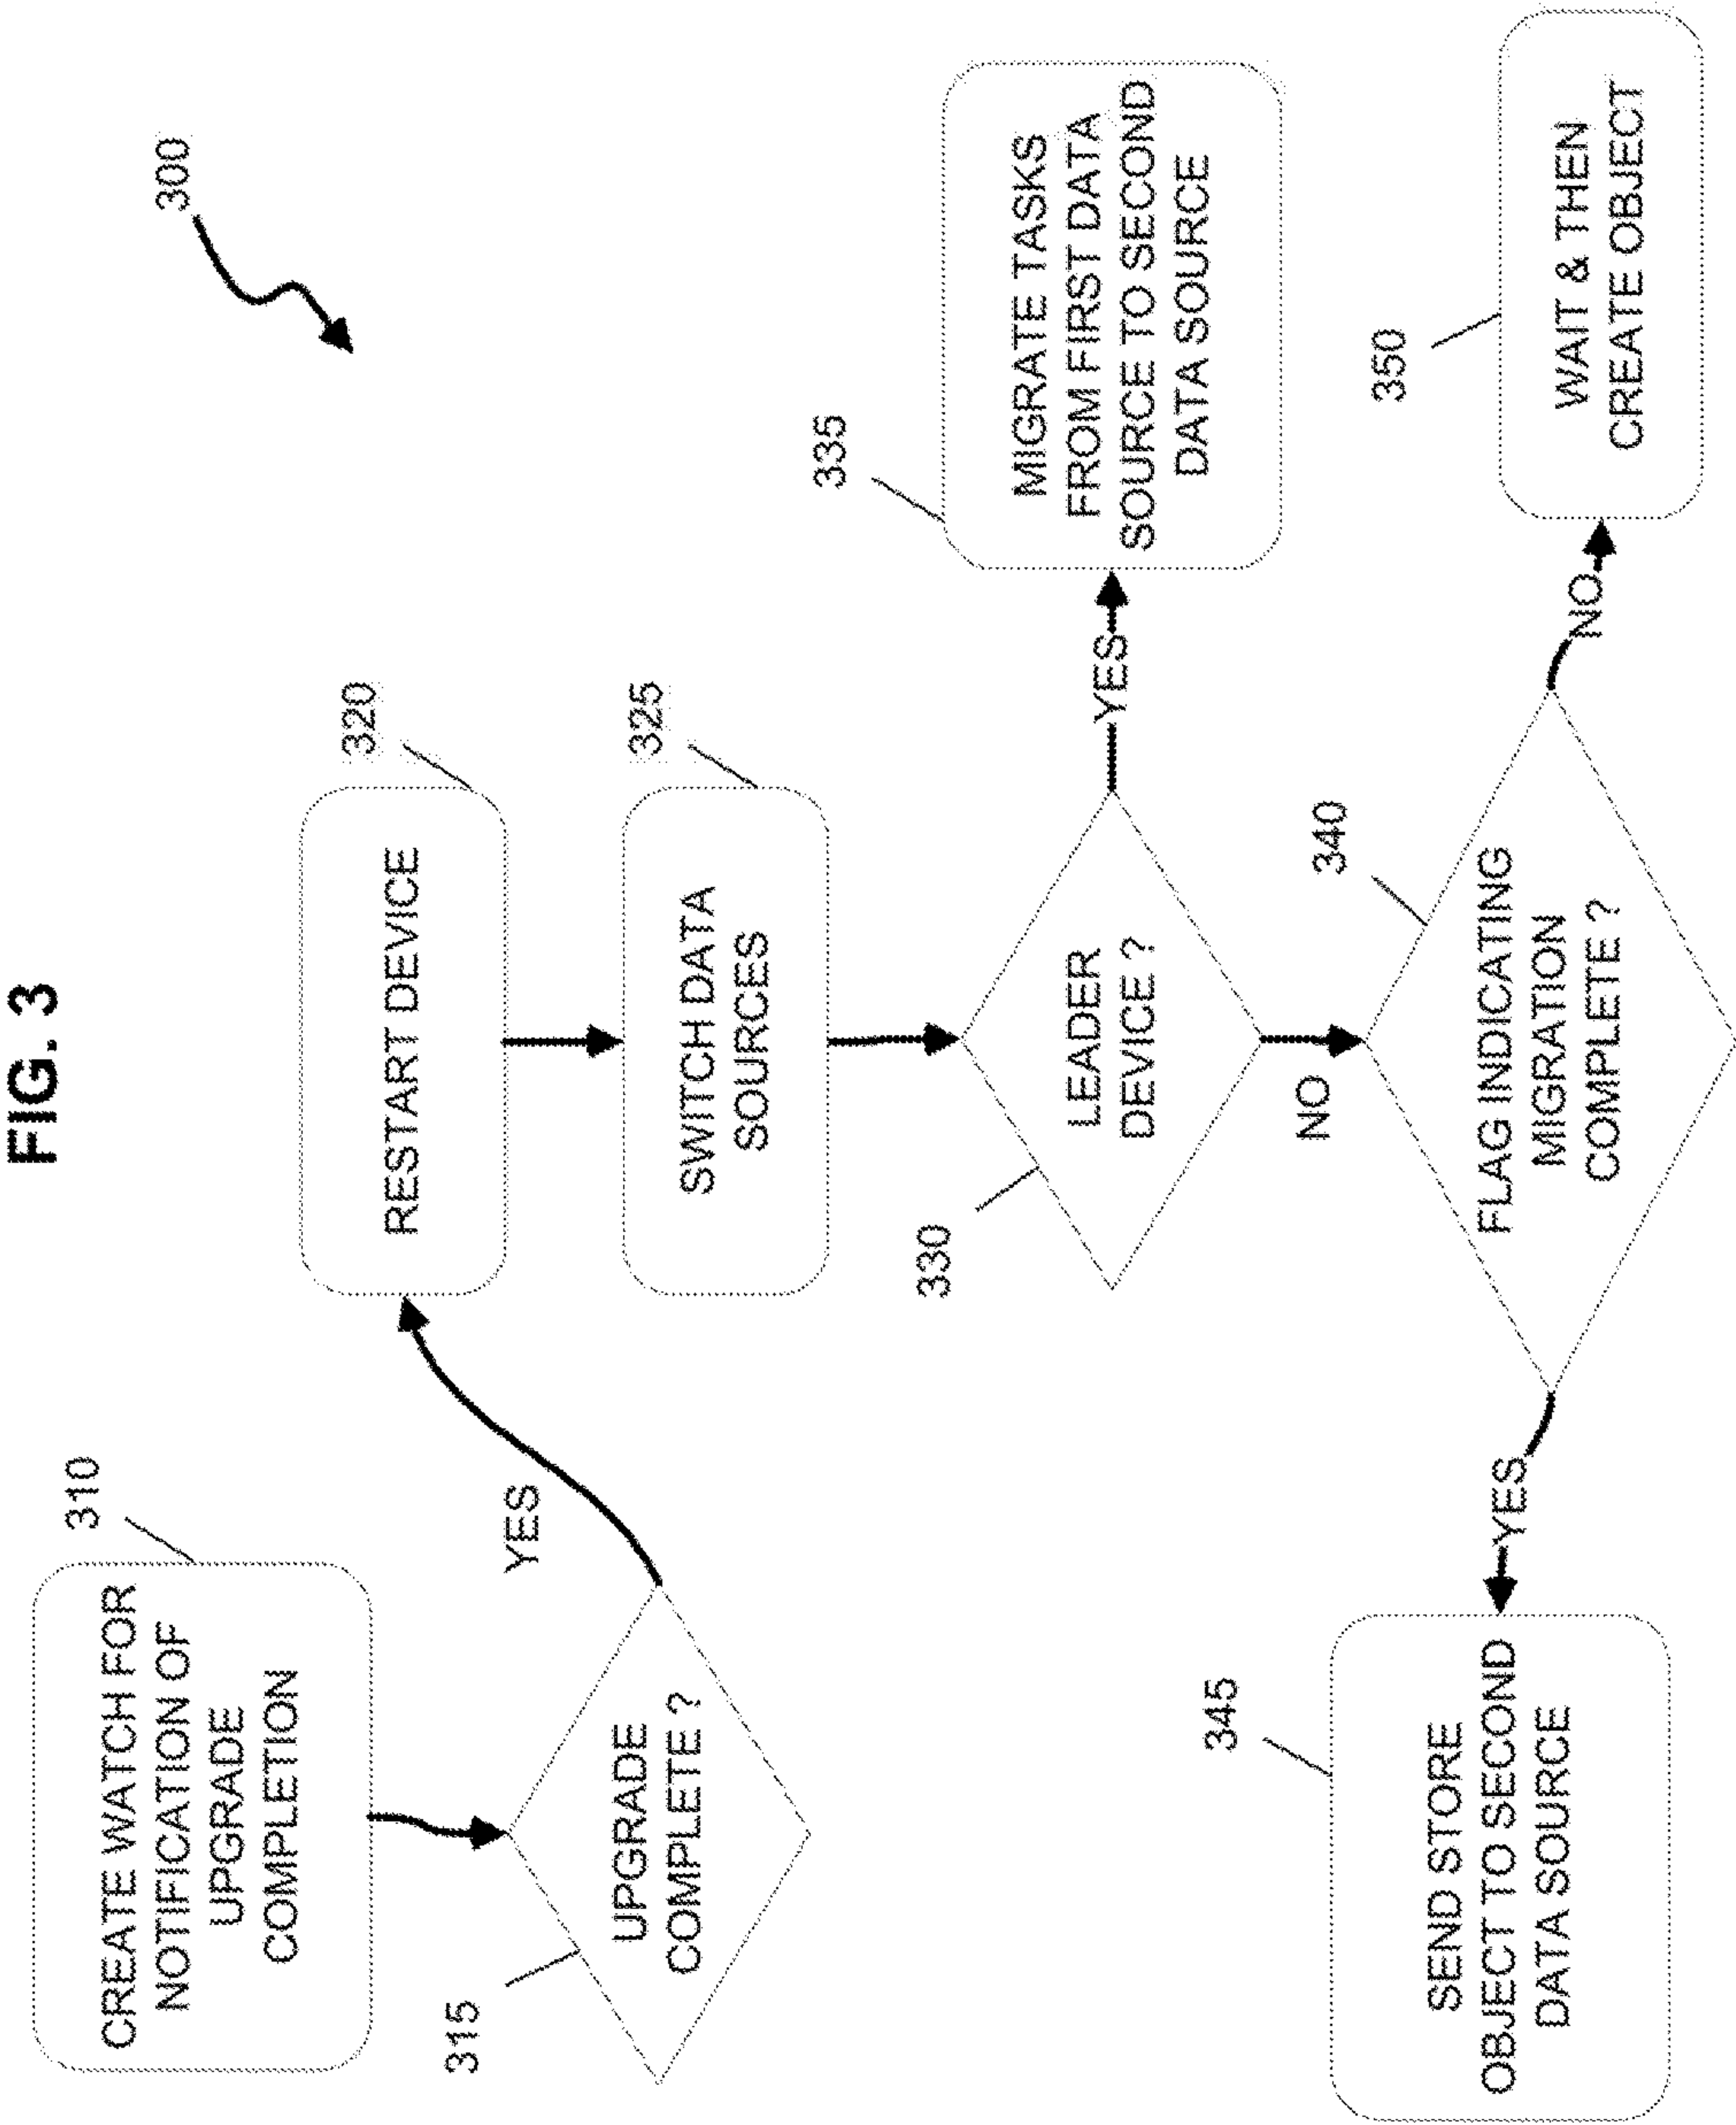


FIG. 2







1

# SWITCHING DATA SOURCES DURING ROLLING UPGRADE OF A LEADER-BASED DISTRIBUTED SERVICE WITH MINIMAL DISRUPTIONS

## CROSS REFERENCE TO RELATED APPLICATION

This application claims priority to U.S. Provisional Application 62/434,873, filed Dec. 15, 2016, incorporated by reference in its entirety.

## BACKGROUND

A virtual machine (“VM”) may refer to a specific software-based implementation of a machine in a virtualization environment, in which the hardware resources of a real computer (e.g., CPU, memory, etc.) are virtualized or transformed into the underlying support for the fully functional virtual machine that can run its own operating system and applications on the underlying physical resources just like a real computer.

## SUMMARY

Aspects of the present disclosure relate generally to a virtualization environment, and more particularly to a system for switching data sources during rolling upgrade of a leader-based distributed service with minimal disruptions.

One implementation disclosed herein is a method of switching data sources during a rolling upgrade of a leader-based distributed service. In at least one implementation, the method includes establishing, at a computing device in a cluster of devices in a network, a watch for notification of completion of a rolling upgrade on all devices in the cluster; restarting the computing device on completion of the rolling upgrade; switching, at the computing device, from a first data source to a second data source once the rolling upgrade is complete on all devices in the cluster; if the computing device is a leader device in the cluster, migrating tasks from the first data source to the second data source after completion of the rolling upgrade; and if the computing device is a non-leader device in the cluster, determining if task migration to the second data source is done and, when done, sending a store object to the second data source. Tasks can include a user activity such as creating a virtual machine, powering on a virtual machine, etc.

In some implementations, the method further includes setting a flag by the leader device to indicate migration is complete. The determining, at the non-leader device, if task migration to the second data can be done based on the set flag. If the flag is not set, the method can include the non-leader device waiting and then proceeding with object creation.

In some implementations, the method further includes the migrated tasks including both pending and recent tasks. The method can also include dispatching tasks from the second data source after the tasks are migrated to the second data source. Further, the rolling upgrade can include changing versions of software running on devices in the cluster of devices.

In another aspect, the present disclosure is directed to a system having a cluster of devices connected in a network. A computing device in the cluster of devices is configured to establish a watch for notification of completion of a rolling upgrade on all devices in the cluster; restart on completion of the rolling upgrade at the computing device; switch from

2

a first data source to a second data source once the rolling upgrade is complete on all devices in the cluster; migrate tasks from the first data source to the second data source after completion of the rolling upgrade if the computing device is a leader device in the cluster; and determine if task migration to the second data source is done if the computing device is a non-leader device in the cluster, and, when done, send a store object to the second data source.

In some implementations, a non-transitory computer readable storage device has a processor-executable instructions embodied thereon, for switching data sources during a rolling upgrade of a leader-based distributed service, where the processor-executable instructions cause a computing device in a cluster of devices in a network to perform operations including: establishing a watch for notification of completion of a rolling upgrade on all devices in the cluster; restarting the computing device on completion of the rolling upgrade at the computing device; switching from a first data source to a second data source once the rolling upgrade is complete on all devices in the cluster; if the computing device is a leader device in the cluster, migrating tasks from the first data source to the second data source after completion of the rolling upgrade; and if the computing device is a non-leader device in the cluster, determining if task migration to the second data source is done and, when done, sending a store object to the second data source.

## BRIEF DESCRIPTION OF THE DRAWINGS

Illustrative embodiments will hereafter be described with reference to the accompanying drawings.

FIG. 1A illustrates a clustered virtualization environment **100** according to particular embodiments.

FIG. 1B illustrates data flow within an example clustered virtualization environment **100** according to particular embodiments.

FIG. 2 is a block diagram of an illustrative computing system **200** suitable for implementing particular embodiments.

FIG. 3 is a flow diagram of an illustrative method **300** suitable for implementing particular embodiments.

Like reference numbers and designations in the various drawings indicate like elements.

## DETAILED DESCRIPTION

Virtualization works by inserting a thin layer of software directly on the computer hardware or on a host operating system. This layer of software contains a virtual machine monitor or “hypervisor” that allocates hardware resources dynamically and transparently. Multiple operating systems run concurrently on a single physical computer and share hardware resources with each other. By encapsulating an entire machine, including CPU, memory, operating system, and network devices, a virtual machine is completely compatible with most standard operating systems, applications, and device drivers. Most modern implementations allow several operating systems and applications to safely run at the same time on a single computer, with each having access to the resources it needs when it needs them.

Virtualization allows one to run multiple virtual machines on a single physical machine, with each virtual machine sharing the resources of that one physical computer across multiple environments. Different virtual machines can run different operating systems and multiple applications on the same physical computer.



One reason for the broad adoption of virtualization in modern business and computing environments is because of the resource utilization advantages provided by virtual machines. Without virtualization, if a physical machine is limited to a single dedicated operating system, then during periods of inactivity by the dedicated operating system the physical machine is not utilized to perform useful work. This is wasteful and inefficient if there are users on other physical machines which are currently waiting for computing resources. To address this problem, virtualization allows multiple VMs to share the underlying physical resources so that during periods of inactivity by one VM, other VMs can take advantage of the resource availability to process workloads. This can produce great efficiencies for the utilization of physical devices, and can result in reduced redundancies and better resource cost management.

Furthermore, there are now products that can aggregate multiple physical machines, running virtualization environments to not only utilize the processing power of the physical devices to aggregate the storage of the individual physical devices to create a logical storage pool wherein the data may be distributed across the physical devices but appears to the virtual machines to be part of the system that the virtual machine is hosted on. Such systems operate under the covers by using metadata, which may be distributed and replicated any number of times across the system, to locate the indicated data. These systems are commonly referred to as clustered systems, wherein the resources of the group are pooled to provide logically combined, but physically separate systems.

Further details of aspects, objects, and advantages of the invention are described below in the detailed description, drawings, and claims. Both the foregoing general description and the following detailed description are exemplary and explanatory, and are not intended to be limiting as to the scope of the invention. Particular embodiments may include all, some, or none of the components, elements, features, functions, operations, or steps of the embodiments disclosed above. The subject matter which can be claimed comprises not only the combinations of features as set out in the attached claims but also any other combination of features in the claims, wherein each feature mentioned in the claims can be combined with any other feature or combination of other features in the claims. Furthermore, any of the embodiments and features described or depicted herein can be claimed in a separate claim and/or in any combination with any embodiment or feature described or depicted herein or with any of the features of the attached claims.

FIG. 1A illustrates a clustered virtualization environment **100** according to particular embodiments. The architecture of FIG. 1A can be implemented for a distributed platform that contains multiple host machines **101a-c** that manage multiple tiers of storage. The multiple tiers of storage may include storage that is accessible through network **140**, such as, by way of example and not limitation, cloud storage **126** (e.g., which may be accessible through the Internet), network-attached storage (NAS) **128** (e.g., which may be accessible through a LAN), or a storage area network (SAN). Unlike the prior art, the present embodiment also permits local storage **122a-c** that is incorporated into or directly attached to the host machine and/or appliance to be managed as part of storage pool **160**. Examples of such local storage include Solid State Drives **125** (henceforth “SSDs”), Hard Disk Drives **127** (henceforth “HDDs” or “spindle drives”), optical disk drives, external drives (e.g., a storage device connected to a host machine via a native drive interface or a serial attached SCSI interface), or any other

direct-attached storage. These storage devices, both direct-attached and network-accessible, collectively form storage pool **160**. Virtual disks (or “vDisks”) may be structured from the physical storage devices in storage pool **160**, as described in more detail below. As used herein, the term vDisk refers to the storage abstraction that is exposed by a Controller/Service VM (CVM) **110** to be used by a user VM **105**. In particular embodiments, the vDisk may be exposed via iSCSI (“internet small computer system interface”) or NFS (“network file system”) and is mounted as a virtual disk on the user VM. In particular embodiments, vDisks may be organized into one or more volume groups (VGs).

Each host machine **101a-c** may run virtualization software, such as VMWARE ESX(I), MICROSOFT HYPER-V, or REDHAT KVM. The virtualization software includes hypervisor **130a-c** to create, manage, and destroy user VMs **105**, as well as managing the interactions between the underlying hardware and user VMs **105**. User VMs **105** may run one or more applications that may operate as “clients” with respect to other elements within virtualization environment **100**. In particular embodiments, one or more lightweight application containers (e.g., DOCKER containers) virtualizing the operating system by encapsulating dependencies (e.g., executable binaries or libraries) within may be deployed on a user VM **105**. Though not depicted in FIG. 1A, a hypervisor may connect to network **140**. In particular embodiments, a host machine **101** may be a physical hardware computing device; in particular embodiments, a host machine **101** may be a virtual machine.

CVMs **110a-c** are used to manage storage and input/output (“I/O”) activities according to particular embodiments. These special VMs act as the storage controller in the currently described architecture. Multiple such storage controllers may coordinate within a cluster to form a unified storage controller system. CVMs **110** may run as virtual machines on the various host machines **101**, and work together to form a distributed system **110** that manages all the storage resources, including local storage **122**, NAS **128**, and cloud storage **126**. The CVMs may connect to network **140** directly, or via a hypervisor. Since the CVMs run independent of hypervisors **130a-c**, this means that the current approach can be used and implemented within any virtual machine architecture, since the CVMs of particular embodiments can be used in conjunction with any hypervisor from any virtualization vendor.

A host machine may be designated as a leader node within a cluster of host machines. For example, host machine **101b**, as indicated by the asterisks, may be a leader node. A leader node may have a software component designated to perform operations of the leader. For example, CVM **110b** on host machine **101b** may be designated to perform such operations. A leader may be responsible for monitoring or handling requests from other host machines or software components on other host machines throughout the virtualized environment. If a leader fails, a new leader may be designated. In particular embodiments, a management module (e.g., in the form of an agent) may be running on the leader node.

Each CVM **110a-c** exports one or more block devices or NFS server targets that appear as disks to user VMs **105a-c**. These disks are virtual, since they are implemented by the software running inside CVMs **110a-c**. Thus, to user VMs **105a-c**, CVMs **110a-c** appear to be exporting a clustered storage appliance that contains some disks. All user data (including the operating system) in the user VMs **105a-c** and reside on these virtual disks.



## 5

Significant performance advantages can be gained by allowing the virtualization system to access and utilize local storage **122** as disclosed herein. This is because I/O performance is typically much faster when performing access to local storage **122** as compared to performing access to NAS **128** across a network **140**. This faster performance for locally attached storage **122** can be increased even further by using certain types of optimized local storage devices, such as SSDs. Further details regarding methods and mechanisms for implementing the virtualization environment illustrated in FIG. 1A are described in U.S. Pat. No. 8,601,473, which is hereby incorporated by reference in its entirety.

FIG. 1B illustrates data flow within an example clustered virtualization environment **100** according to particular embodiments. As described above, one or more user VMs and a CVM may run on each host machine **101** along with a hypervisor. As a user VM performs I/O operations (e.g., a read operation or a write operation), the I/O commands of the user VM may be sent to the hypervisor that shares the same server as the user VM. For example, the hypervisor may present to the virtual machines an emulated storage controller, receive an I/O command and facilitate the performance of the I/O command (e.g., via interfacing with storage that is the object of the command, or passing the command to a service that will perform the I/O command). An emulated storage controller may facilitate I/O operations between a user VM and a vDisk. A vDisk may present to a user VM as one or more discrete storage drives, but each vDisk may correspond to any part of one or more drives within storage pool **160**. Additionally or alternatively, CVM **110a-c** may present an emulated storage controller either to the hypervisor or to user VMs to facilitate I/O operations. CVM **110a-c** may be connected to storage within storage pool **160**. CVM **110a** may have the ability to perform I/O operations using local storage **122a** within the same host machine **101a**, by connecting via network **140** to cloud storage **126** or NAS **128**, or by connecting via network **140** to local storage **122b-c** within another host machine **101b-c** (e.g., via connecting to another CVM **110b** or **110c**). In particular embodiments, any suitable computing system **200** may be used to implement a host machine **101**.

FIG. 2 is a block diagram of an illustrative computing system **200** suitable for implementing particular embodiments. In particular embodiments, one or more computer systems **200** perform one or more steps of one or more methods described or illustrated herein. In particular embodiments, one or more computer systems **200** provide functionality described or illustrated herein. In particular embodiments, software running on one or more computer systems **200** performs one or more steps of one or more methods described or illustrated herein or provides functionality described or illustrated herein. Particular embodiments include one or more portions of one or more computer systems **200**. Herein, reference to a computer system may encompass a computing device, and vice versa, where appropriate. Moreover, reference to a computer system may encompass one or more computer systems, where appropriate.

This disclosure contemplates any suitable number of computer systems **200**. This disclosure contemplates computer system **200** taking any suitable physical form. As example and not by way of limitation, computer system **200** may be an embedded computer system, a system-on-chip (SOC), a single-board computer system (SBC) (such as, for example, a computer-on-module (COM) or system-on-module (SOM)), a desktop computer system, a mainframe, a mesh of computer systems, a server, a laptop or notebook

## 6

computer system, a tablet computer system, or a combination of two or more of these. Where appropriate, computer system **200** may include one or more computer systems **200**; be unitary or distributed; span multiple locations; span multiple machines; span multiple data centers; or reside in a cloud, which may include one or more cloud components in one or more networks. Where appropriate, one or more computer systems **200** may perform without substantial spatial or temporal limitation one or more steps of one or more methods described or illustrated herein. As an example and not by way of limitation, one or more computer systems **200** may perform in real time or in batch mode one or more steps of one or more methods described or illustrated herein. One or more computer systems **200** may perform at different times or at different locations one or more steps of one or more methods described or illustrated herein, where appropriate.

Computer system **200** includes a bus **202** (e.g., an address bus and a data bus) or other communication mechanism for communicating information, which interconnects subsystems and devices, such as processor **204**, memory **206** (e.g., RAM), static storage **208** (e.g., ROM), dynamic storage **210** (e.g., magnetic or optical), communication interface **214** (e.g., modem, Ethernet card, a network interface controller (NIC) or network adapter for communicating with an Ethernet or other wire-based network, a wireless NIC (WNIC) or wireless adapter for communicating with a wireless network, such as a WI-FI network), input/output (I/O) interface **212** (e.g., keyboard, keypad, mouse, microphone). In particular embodiments, computer system **200** may include one or more of any such components.

In particular embodiments, processor **204** includes hardware for executing instructions, such as those making up a computer program. As an example and not by way of limitation, to execute instructions, processor **204** may retrieve (or fetch) the instructions from an internal register, an internal cache, memory **206**, static storage **208**, or dynamic storage **210**; decode and execute them; and then write one or more results to an internal register, an internal cache, memory **206**, static storage **208**, or dynamic storage **210**. In particular embodiments, processor **204** may include one or more internal caches for data, instructions, or addresses. This disclosure contemplates processor **204** including any suitable number of any suitable internal caches, where appropriate. As an example and not by way of limitation, processor **204** may include one or more instruction caches, one or more data caches, and one or more translation lookaside buffers (TLBs). Instructions in the instruction caches may be copies of instructions in memory **206**, static storage **208**, or dynamic storage **210**, and the instruction caches may speed up retrieval of those instructions by processor **204**. Data in the data caches may be copies of data in memory **206**, static storage **208**, or dynamic storage **210** for instructions executing at processor **204** to operate on; the results of previous instructions executed at processor **204** for access by subsequent instructions executing at processor **204** or for writing to memory **206**, static storage **208**, or dynamic storage **210**; or other suitable data. The data caches may speed up read or write operations by processor **204**. The TLBs may speed up virtual-address translation for processor **204**. In particular embodiments, processor **204** may include one or more internal registers for data, instructions, or addresses. This disclosure contemplates processor **204** including any suitable number of any suitable internal registers, where appropriate. Where appropriate, processor **204** may include one or more arithmetic logic units (ALUs); be a multi-core processor; or include



one or more processors **202**. Although this disclosure describes and illustrates a particular processor, this disclosure contemplates any suitable processor.

In particular embodiments, I/O interface **212** includes hardware, software, or both, providing one or more interfaces for communication between computer system **200** and one or more I/O devices. Computer system **200** may include one or more of these I/O devices, where appropriate. One or more of these I/O devices may enable communication between a person and computer system **200**. As an example and not by way of limitation, an I/O device may include a keyboard, keypad, microphone, monitor, mouse, printer, scanner, speaker, still camera, stylus, tablet, touch screen, trackball, video camera, another suitable I/O device or a combination of two or more of these. An I/O device may include one or more sensors. This disclosure contemplates any suitable I/O devices and any suitable I/O interfaces **212** for them. Where appropriate, I/O interface **212** may include one or more device or software drivers enabling processor **204** to drive one or more of these I/O devices. I/O interface **212** may include one or more I/O interfaces **212**, where appropriate. Although this disclosure describes and illustrates a particular I/O interface, this disclosure contemplates any suitable I/O interface.

In particular embodiments, communication interface **214** includes hardware, software, or both providing one or more interfaces for communication (such as, for example, packet-based communication) between computer system **200** and one or more other computer systems **200** or one or more networks. As an example and not by way of limitation, communication interface **214** may include a network interface controller (NIC) or network adapter for communicating with an Ethernet or other wire-based network or a wireless NIC (WNIC) or wireless adapter for communicating with a wireless network, such as a WI-FI network. This disclosure contemplates any suitable network and any suitable communication interface **214** for it. As an example and not by way of limitation, computer system **200** may communicate with an ad hoc network, a personal area network (PAN), a local area network (LAN), a wide area network (WAN), a metropolitan area network (MAN), or one or more portions of the Internet or a combination of two or more of these. One or more portions of one or more of these networks may be wired or wireless. As an example, computer system **200** may communicate with a wireless PAN (WPAN) (such as, for example, a BLUETOOTH WPAN), a WI-FI network, a WI-MAX network, a cellular telephone network (such as, for example, a Global System for Mobile Communications (GSM) network), or other suitable wireless network or a combination of two or more of these. Computer system **200** may include any suitable communication interface **214** for any of these networks, where appropriate. Communication interface **214** may include one or more communication interfaces **214**, where appropriate. Although this disclosure describes and illustrates a particular communication interface, this disclosure contemplates any suitable communication interface.

One or more memory buses (which may each include an address bus and a data bus) may couple processor **204** to memory **206**. Bus **202** may include one or more memory buses, as described below. In particular embodiments, one or more memory management units (MMUs) reside between processor **204** and memory **206** and facilitate accesses to memory **206** requested by processor **204**. In particular embodiments, memory **206** includes random access memory (RAM). This RAM may be volatile memory, where appropriate where appropriate, this RAM may be dynamic RAM

(DRAM) or static RAM (SRAM). Moreover, where appropriate, this RAM may be single-ported or multi-ported RAM. This disclosure contemplates any suitable RAM. Memory **206** may include one or more memories **206**, where appropriate. Although this disclosure describes and illustrates particular memory, this disclosure contemplates any suitable memory.

Where appropriate, the ROM may be mask-programmed ROM, programmable ROM (PROM), erasable PROM (EPROM), electrically erasable PROM (EEPROM), electrically alterable ROM (EAROM), or flash memory or a combination of two or more of these. In particular embodiments, dynamic storage **210** may include a hard disk drive (HDD), a floppy disk drive, flash memory, an optical disc, a magneto-optical disc, magnetic tape, or a Universal Serial Bus (USB) drive or a combination of two or more of these. Dynamic storage **210** may include removable or non-removable (or fixed) media, where appropriate. Dynamic storage **210** may be internal or external to computer system **200**, where appropriate. This disclosure contemplates mass dynamic storage **210** taking any suitable physical form. Dynamic storage **210** may include one or more storage control units facilitating communication between processor **204** and dynamic storage **210**, where appropriate.

In particular embodiments, bus **202** includes hardware, software, or both coupling components of computer system **200** to each other. As an example and not by way of limitation, bus **202** may include an Accelerated Graphics Port (AGP) or other graphics bus, an Enhanced Industry Standard Architecture (EISA) bus, a front-side bus (FSB), a HYPERTRANSPORT (HT) interconnect, an Industry Standard Architecture (ISA) bus, an INFINIBAND interconnect, a low-pin-count (LPC) bus, a memory bus, a Micro Channel Architecture (MCA) bus, a Peripheral Component Interconnect (PCI) bus, a PCI-Express (PCIe) bus, a serial advanced technology attachment (SATA) bus, a Video Electronics Standards Association local (VLB) bus, or another suitable bus or a combination of two or more of these. Bus **202** may include one or more buses **206**, where appropriate. Although this disclosure describes and illustrates a particular bus, this disclosure contemplates any suitable bus or interconnect.

According to particular embodiments, computer system **200** performs specific operations by processor **204** executing one or more sequences of one or more instructions contained in memory **206**. Such instructions may be read into memory **206** from another computer readable/usable medium, such as static storage **208** or dynamic storage **210**. In alternative embodiments, hard-wired circuitry may be used in place of or in combination with software instructions to implement particular embodiments. Thus, particular embodiments are not limited to any specific combination of hardware circuitry and/or software. In particular embodiments, the term “logic” shall mean any combination of software or hardware that is used.

The term “computer readable medium” or “computer usable medium” as used herein refers to any medium that participates in providing instructions to processor **204** for execution. Such a medium may take many forms, including but not limited to, nonvolatile media and volatile media. Non-volatile media includes, for example, optical or magnetic disks, such as static storage **208** or dynamic storage **210**. Volatile media includes dynamic memory, such as memory **206**.

Common forms of computer readable media include, for example, floppy disk, flexible disk, hard disk, magnetic tape, any other magnetic medium, CD-ROM, any other optical medium, punch cards, paper tape, any other physical



medium with patterns of holes, RAM, PROM, EPROM, FLASH-EPROM, any other memory chip or cartridge, or any other medium from which a computer can read.

In particular embodiments, execution of the sequences of instructions may be performed by a single computer system **200**; in alternative embodiments, two or more computer systems **200** coupled by communication link **216** (e.g., LAN, PTSN, or wireless network) may perform the sequence of instructions in coordination with one another.

Computer system **200** may transmit and receive messages, data, and instructions, including program, i.e., application code, through communication link **216** and communication interface **214**. Received program code may be executed by processor **204** as it is received, and/or stored in static storage **208** or dynamic storage **210**, or other non-volatile storage for later execution. A database **220** may be used to store data accessible by the system **200** by way of data interface **218**.

FIG. 3 illustrates a flow diagram of a method **300** showing operations in an illustrative method of switching data sources during a rolling upgrade of a leader-based distributed service. Additional, fewer, or different operations may be performed in method **300** depending on the implementation or embodiment. In an operation **310**, a computing device in a cluster of devices in a network creates a watch for notification of completion of a rolling upgrade on all devices in the cluster. The computing device waits until the watch notifies it that the update is complete across the entire cluster (operation **315**). In an operation **320**, the computing device is restarted upon completion of the rolling upgrade. Once the rolling upgrade is complete on all devices in the cluster and after the restart, the computing device switches from a first data source to a second data source (operation **325**). Pending and recent tasks should only be migrated to a new data source after migration is complete. If the computing device is a leader device for the cluster (operation **330**), the computing device migrates tasks from the first data source to the second data source (operation **335**). If the computing device is a non-leader device in the cluster, the computing device determines if task migration to the second data source is done (operation **340**) and, when done, sends a store object to the second data source (operation **345**). If an indication that the migration is not done, the computing device can wait for a time and then proceed with object creation (operation **350**).

The method enables, for example, rolling updates of an operating system to a cluster of computing devices. In some circumstances, a cluster leader cannot be selected or chosen and, therefore, there is a need to ensure correctness in task migration and execution during updates. Also, method enables all data store switches during a rolling upgrade of leader-based distributed services. Moreover, intent-based tasks are migrated in the correct order.

A system utilizing the method **300** can preserve the history or order of tasks. Tasks maintained in the databases may comprise multiple steps and the database preserves the task status and recovery information, including completed steps and output. The system can track in-progress tasks and tasks to be completed. Advantageously, each node is updated and then each node sets a watch to notice when the software update completes across the entire cluster. Each upgraded node still points to the first database. Once the upgrade across the cluster completes, the nodes are notified and a restart is done. After the restart, a cluster leader is elected and the leader starts migration beginning with the oldest tasks. Once the migration is complete, a flag is set in the key

value store. The first database is retired and a second database starts executing tasks.

Herein, a computer-readable non-transitory storage medium or media may include one or more semiconductor-based or other integrated circuits (ICs) (such, as for example, field-programmable gate arrays (FPGAs) or application-specific ICs (ASICs)), hard disk drives (HDDs), hybrid hard drives (HHDs), optical discs, optical disc drives (ODDs), magneto-optical discs, magneto-optical drives, floppy diskettes, floppy disk drives (FDDs), magnetic tapes, solid-state drives (SSDs), RAM-drives, SECURE DIGITAL cards or drives, any other suitable computer-readable non-transitory storage media, or any suitable combination of two or more of these, where appropriate. A computer-readable non-transitory storage medium may be volatile, non-volatile, or a combination of volatile and non-volatile, where appropriate.

Herein, “or” is inclusive and not exclusive, unless expressly indicated otherwise or indicated otherwise by context. Therefore, herein, “A or B” means “A, B, or both,” unless expressly indicated otherwise or indicated otherwise by context. Moreover, “and” is both joint and several, unless expressly indicated otherwise or indicated otherwise by context. Therefore, herein, “A and B” means “A and B, jointly or severally,” unless expressly indicated otherwise or indicated otherwise by context.

The scope of this disclosure encompasses all changes, substitutions, variations, alterations, and modifications to the example embodiments described or illustrated herein that a person having ordinary skill in the art would comprehend. The scope of this disclosure is not limited to the example embodiments described or illustrated herein. Moreover, although this disclosure describes and illustrates respective embodiments herein as including particular components, elements, feature, functions, operations, or steps, any of these embodiments may include any combination or permutation of any of the components, elements, features, functions, operations, or steps described or illustrated anywhere herein that a person having ordinary skill in the art would comprehend. Furthermore, reference in the appended claims to an apparatus or system or a component of an apparatus or system being adapted to, arranged to, capable of, configured to, enabled to, operable to, or operative to perform a particular function encompasses that apparatus, system, component, whether or not it or that particular function is activated, turned on, or unlocked, as long as that apparatus, system, or component is so adapted, arranged, capable, configured, enabled, operable, or operative.

What is claimed is:

1. A method comprising:

determining, by each of a plurality of computing devices of a cluster, that an upgrade is complete on all of the plurality of computing devices of the cluster;  
in response to the determining, restarting each of the plurality of computing devices upon completion of the upgrade on all of the plurality of computing devices;  
and

upon restarting of all of the plurality of computing devices:

designating one of the plurality of computing devices as a leader device;  
switching each of the plurality of computing devices from a first data source to a second data source;  
migrating, by the leader device, a task in the first data source to the second data source;  
in response to the determining, by a non-leader device that the migration of tasks is complete, executing, by



## 11

the non-leader device of the plurality of computing devices, the task from the second data source instead of the first data source upon migration of the task to the second data source.

2. The method of claim 1, further comprising setting a flag by the leader device to indicate that the migration of the task is complete.

3. The method of claim 2, further comprising determining, by the non-leader device, that the migration of the task is complete based on the set flag.

4. The method of claim 2, further comprising waiting, by the non-leader device, a predetermined period of time for the flag to be set before executing the task from the second data source.

5. The method of claim 1, wherein the task that is migrated comprises a pending task or a recent task.

6. The method of claim 1 further comprising establishing, by each of the plurality of computing devices, a watch for determining that the upgrade is complete on all of the plurality of computing devices of the cluster.

7. The method of claim 5, wherein each of the plurality of computing devices establishes the watch upon completing the upgrade.

8. The method of claim 1, further comprising additional tasks, and wherein the leader device migrates the task and the additional tasks to the second data source from the first data source in a predetermined order.

9. A system comprising:

a plurality of computing devices of a cluster, each of the plurality of computing devices comprising a processor and a memory communicatively coupled to the processor, the memory storing computer-readable instructions that, when executed by the processor, cause the processor to:

determine that an upgrade is complete on all of the plurality of computing devices of the cluster;

in response to determining the upgrade is complete, restart all of the plurality of computing devices; and upon restarting of all of the plurality of computing devices:

one of the plurality of computing devices is designated as a leader device;

switch each of the plurality of computing devices from a first data source to a second data source upon restarting;

the leader device migrates a task in the first data source to the second data source; and

a non-leader device of the plurality of computing devices executes the task from the second data source instead of the first data source upon determining migration of the task to the second data source is complete.

10. The system of claim 9, wherein the leader device sets a flag to indicate that the migration of the task is complete.

11. The system of claim 10, wherein the non-leader device determines that the migration of the task is complete based on the set flag.

## 12

12. The system of claim 10, wherein the non-leader device waits for a predetermined period of time for the flag to be set before executing the task from the second data source.

13. The system of claim 9, further comprising additional tasks, and wherein the leader device migrates the task and the additional tasks to the second data source from the first data source in a predetermined order.

14. The system of claim 9, wherein each of the plurality of computing devices comprises computer-readable instructions to establish a watch to determine that the upgrade is complete on all of the plurality of computing devices of the cluster.

15. A non-transitory computer-readable media comprising computer-executable instructions embodied thereon that, when executed by a processor of each of a computing device of a plurality of computing devices of a cluster cause the respective processor to:

determine that an upgrade is complete on all of the plurality of computing devices of the cluster;

in response to the determining, restart all of the plurality of computing devices; and

upon restarting of all of the plurality of computing devices:

one of the plurality of computing devices is designated as a leader device;

switch each of the plurality of computing devices from a first data source to a second data source,

the leader device migrates a task in the first data source to the second data source; and

a non-leader device of the plurality of computing devices executes the task from the second data source instead of the first data source upon determining migration of the task to the second data source is complete.

16. The non-transitory computer-readable media of claim 15, wherein the leader device sets a flag to indicate that the migration of the task is complete.

17. The non-transitory computer-readable media of claim 16, wherein the non-leader device determines that the migration of the task is complete based on the set flag.

18. The non-transitory computer-readable media of claim 16, wherein the non-leader device waits for a predetermined period of time for the flag to be set before executing the task from the second data source.

19. The non-transitory computer-readable media of claim 15, further comprising additional tasks, and wherein the leader device migrates the task and the additional tasks to the second data source from the first data source in a predetermined order.

20. The non-transitory computer-readable media of claim 15, wherein the processor further comprises computer-readable instructions to establish a watch to determine that the upgrade is complete on all of the plurality of computing devices of the cluster.

21. The non-transitory computer-readable media of claim 20, wherein the processor establishes the watch upon completing the upgrade.

\* \* \* \* \*